PCT/GB 00/00982
09/936287

Eur päisches Patentamt
Eur pean Patent Office
Office eur péen des brevets

GB 00/982

REC'D 28 JUN 2000
WIPO PCT

Bescheinigung Certificate Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

99305451.9

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE, 22/05/00
LA HAYE, LE

THIS PAGE BLANK (USPTO)

Europäisches Patentamt
Eur pean Patent Office
Office eur péen des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.: 99305451.9
Demande n°:

Anmeldetag:
Date of filing: 08/07/99
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
BRITISH TELECOMMUNICATIONS public limited company
London EC1A 7AJ
UNITED KINGDOM

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:
Progressive routing in a communications network

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat: State: Pays:	Tag: Date: Date:	Aktenzeichen: File no. Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:
H04Q3/66, H04Q11/04, H04L12/56

Am Anmeldetag benannte Vertragstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE
Etats contractants désignés lors du depôt:

Bemerkungen:
Remarks:
Remarques:
See for original title of the application page 1 of the description.

THIS PAGE BLANK (USPTO)

ROUTING IN A COMMUNICATIONS NETWORK

This invention relates to a method of routing in a communications network of interconnected nodes, and particularly, but not exclusively, in a sparsely connected network.

A number of routing algorithms are known for routing in a network of interconnected nodes. For example, in the event of a fault preventing a message from being forwarded from a transit node T1 to an adjacent node T2, the message is sent on an alternative route to T2 via another transit node T3. In another example, if there is a fault on a primary route to a destination node, the message is returned to the source node and a secondary route is tried from the source node to the destination node.

In accordance with one aspect of the present invention, there is provided a method of routing in a communications network of interconnected nodes, the nodes being arranged to generate messages, each message having a destination information element containing the identity of a destination node for that message, a source information element containing the identity of the source node of that message, and a virtual source information element initially containing the identity of that source node, the method comprising performing at each node the steps of:

(a) retrieving from a message the contents of its destination information element and its virtual source information element;

(b) comparing the retrieved destination node identity with its own node identity; and, when it is not the destination node for that message,

(c) comparing the retrieved virtual source node identity with its own node identity and,

in the event of a match at step (c),

(d) sending the message to the destination node via the highest ranking available route of a ranked set of alternative routes for that destination node, and

in the event of a mismatch at step (c),

(e) sending the message to the destination node via a predetermined alternative route for that destination node; and

in the event of the unavailability of the predetermined alternativ route in step (e),

(f) replacing the content of the virtual source information element of the message with its own node identity and performing step (d), and

if no route is available,

(g) replacing the content of the virtual source information element of the message with the node identity of the node from which the message was received and sending the message back to that node from which it was received.

Preferably, at least two of the set of alternative routes are node-disjoint routes.

Preferably step (g) further comprises changing the state of a flag in a crankback information element of the message, and step (d) comprises ignoring the alternative route that had previously been used for that message.

In accordance with another aspect of the present invention, there is provided a node for use in a communications network of interconnected nodes, the node being arranged to generate messages, each message having a destination information element containing the identity of a destination node for that message, a source information element containing the identity of the source node of that message, and a virtual source information element initially containing the identity of that source node, and being arranged

to retrieve from a message the contents of its destination information element and its virtual source information element;

to compare the retrieved destination node identity with its own nod identity; and, when it is not the destination node for that message,

to compare the retrieved virtual source node identity with its own node identity and,

in the event of a match between the retrieved virtual source node identity and its own node identity,

to send the message to the destination node via the highest ranking available route of a ranked set of alternative routes for that destination node, and

in the event of a mismatch between the retrieved virtual source node identity and its own node identity,

to send the message to the destination node via a predetermined alternative route for that destination node; and

in the event of the unavailability of said predetermined alternative route,

to replace the content of the virtual source information element of the message with its own node identity and to send the message to the destination node via the highest ranking available route of the ranked set of alternative routes for that destination node, and

if no route is available,

to replace the content of the virtual source information element of the message with the node identity of the node from which the message was received and to send the message back to that node from which it was received.

In accordance with a further aspect of the present invention, there is provided a communications network of interconnected nodes, each of the nodes being as defined in the preceding paragraph.

A specific embodiment of a method in accordance with the present invention will now be described by way of example with reference to the drawings, in which:-

Figure 1 shows part of a sparsely connected network;

Figure 2 shows information elements of a message; and

Figures 3 to 6 respectively show routing tables of some of the nodes of the network of Figure 1.

Before proceeding to the detailed description, the reader may find it useful to have definitions of some of the terms in this art.

Crankback refers to a mechanism for re-routing circuits which have either been broken due to the failure of some network element, or else have been unable to be established along their designated routes because of a change in network conditions since the 'topology state database' from which the routes were computed was last updated.

Crankback to source is when a call arrives at a switch (i.e. node) and it cannot be forwarded to the next switch designated in its designated transit list (DTL) or other route indicator (referred to as a routing table herein), a message is sent to the originating switch of the DTL or the call, requiring the call to be re-routed on a separate route.

Hop by hop crankback is when a call arrives at a switch and it cannot be forwarded to the next stage on its rout , a message is sent t the previous switch

on the route requiring the call to be re-routed in such a way as to avoid the switch where it previously stalled.

Limited loop prevention is where, if a switch attempts to route a call setup request (message) back to the switch from which it has just received that call setup, i.e. attempts to perform a "u-turn", then this condition will be recognised and the switch will be prevented from sending the request to that switch.

In Figure 1, a network 10 comprises a multiplicity of switching nodes NX, where X is a node identifier, and interconnecting links LXY, where X and Y are terminating node identifiers for that link. As an example, the link interconnecting nodes NS and NA is arbitrarily designated LSA, although it could equally be designated LAS,

The nodes NX are arranged to switch traffic being carried in accordance with international standards for asynchronous transfer mode (ATM), and although, for convenience, only ten nodes are shown, in a practical network, there will be many more nodes, e.g. in the planned UK ATM network there will be about 100 nodes. The present invention is not limited to ATM networks, thus in variants th nodes can be arranged for switching traffic being carried in accordance with other standards, e.g. plesiochronous digital hierarchy or synchronous digital hierarchy using CCITT No 7 signalling system, and packet switching systems.

The network 10 is partially meshed, in other words, not every node NX is connected to every other node NX. If the network were fully meshed, also kn wn as a fully connected, or fully interconnected network, there would be n(n-1)/2 links LXY where n is the total number of nodes in the network, but in situations where the present invention is particularly advantageous, the network 10 has considerably fewer links LXY, and such a network is referred to as a sparsely connected network. Typically, a sparsely connected network has less than half the number of links LXY of a fully meshed network,.

To illustrate the routing method of the present invention, one of the nodes is designated as a source node NS, another node is designated as a destination node ND, the other nodes in Figure 1 are designated NA, NB, NC, NE, NF, NG, NH and NJ.

In a sparsely connected network, each of the nod s stores, for use in routing messages for which it is the actual or the virtual source, a respective set

of ranked alternative routes, comprising a respective primary pre-planned route and at least a secondary planned route, to each other of the nodes. As described in more detail below, the primary, i.e. highest ranking, route is to be tried first for calls for which the node is the actual source or the virtual source, and, when the primary route is not available, e.g. because of a link failure or a node failure, the next highest ranking route is tried, and so on, depending upon the number of alternative routes in the set.

In this embodiment, the routes in each respective set are node-disjoint routes, in other words, other than the source and destination nodes, they do not have any other node in common. However, in some sparsely connected networks it may not be possible or desirable for all the routes in a set to be node-disjoint routes, but the present invention will still work advantageously.

Suppose that there is a new call at node NS for node ND, and that the primary route is via link LSA to node NA, link LAB to node NB, link LBH to node NH, and finally link LHD to node ND and the secondary route is via link LSE to node NE, link LEF to node NF, link LFG to node NG, and finally link LGD to node ND.

The node NS will generate a Setup Request message 30, also known as a Routing Request, shown in Figure 2, comprising the known information elements 32 of the standard ATM Setup Request message, e.g. for source node identity 32S, destination node identity 32D, and an information element for data 38. The message 30 includes an additional information element 34, which will be referred to as the virtual source information element, and an additional information element 36 containing a crankback flag whose normal state is reset. When a node acts as source node and generates the Setup Request message 30, it will insert its own identity in the normal information element 32S for source node identity and also in the virtual source information element 34.

Each node NX has a respective routing table 20 (e.g. routing table 20S shown in Figure 3) comprising a first column for the identity of the virtual source node. In practice, the node will retrieve the contents of the virtual source information element 34 of any Setup Request message 30 that it handles, and refer to the routing table 20 on the basis of the retrieved contents. This means that when it generates a Setup Request message 30, the contents of the virtual

source information element 34 will initially be its own identity. Otherwise, th virtual source node identity in information element 34 of a received Setup Request message 30 will depend upon the routing history of the received Setup Request message 30. The routing table 20 has a second column for the identity of the destination node, and a third column for the identity of the adjacent node to which the message is to be forwarded. In this embodiment, this third column is referred to as the Address column, and its entries are node identities. In a variant, as is known in the art, another way of identifying the outgoing route is by outgoing link identity.

In each routing table 20 there is a single entry for each source/destination pair in which the source identity is not that of the associated node in the first column for each other node NX, but there are always two entries, a primary route and a secondary route, and possibly one or more further routes, for each destination node for the node NX at which that routing table is resident.

The routing table 20S of source node NS is shown in Figure 3. For the ten node network 10 of Figure 1, there will be nine source/destination pairs (S/X) but only the sets for the source/destination pairs for the nodes ND, NA and NF are shown. In the routing table 20S, the set S/D contains a primary route S1 and a secondary route S2, as does the set S/A, but the set S/F contains an additional tertiary route S3.

For destination node ND, the address for the first entry, primary route S1, is A, and the address for the second entry, secondary route S2, is E. The routing table 20S also contains, for each of the nine destination nodes reachable from node NS, a respective single entry for the other eight nodes for which node NS is a transit node F. For example, for node NF as destination, there are theoretically entries for nodes NA, NB, NC, ND, NE, NG, NH and NJ, but only the last of these is shown in Figure 3.

The source node NS will send the Setup Request message to the nod having the address, i.e. identity, A, associated with S1 in the routing table 20S at the source node NS. Upon receipt of this message, node NA will, in usual manner, retrieve the identity of the destination node from the destination information element 32D and check to see whether the destination node identity match s the node identity NA, i.e. whether node NA is to captur the message for an

associated terminal or whether it is to send the message on to another node in the network. If node NA is not the destination node for that messag , it will then, if it has not already done so, retrieve the identity of the virtual source node from the virtual source information element 34, access its routing table 20A, Figure 4, in respect of the source/destination pair SD, using the retrieved virtual source node identity, find only a single entry, having the address B, and forward the message to node NB.

Upon receipt of this message, node NB will similarly check to see if it is the destination for that message, read the identity of the virtual source node from the virtual source information element 34, access its routing table 20B, Figure 5, in respect of the source/destination pair SD, using the retrieved virtual source node identity, find only a single entry, having the address H, and forward the message to node NH. Node NH will perform the same steps, and similarly forward the message to the destination node ND.

Assuming now that there is a fault, either at the node NH or in the link LBH, and that node NB ascertains by known means, e.g. alarm messages, failure messages or a timeout, that the attempt to forward the Setup Request message to node NH has failed. Node NB now does two things: it replaces, i.e. overwrites, the current, i.e. in this case, initial, contents (S) of the virtual source information element 34 with its own node identity, B; and it accesses its routing table 20B to find the set, in this case, pair, of entries, B1 (H) and B2 (C), for the source/destination pair BD, disregards the route which it now knows is a failure route, i.e. in this case, the route to node NH, and retrieves the address, C, of the route which has not yet been tried. In other words, the primary route from node NB to node ND is, in this case, part of the primary route from node NS to node ND, so the address for the first entry B1 is H, and the address for the second entry B2 is C.

The nodes of network 10 are arranged to prevent a "u-turn", i.e. where a transit node for whatever reason routes the message back along the route from which it was received. In some networks, the nodes may not be so arranged, and in this cas when such a node receives back a message that it has just sent, i.e. the route involves a "u-turn", this, inter alia, constitutes the rout to that next transit node being unavailable.

In variants, it is not always the case that the primary route from node NB to node ND is part of the primary route from node NS to node ND, and it may be that the primary route from node NB to node ND will be via node NC, and the secondary route from node NB to node ND will be via node NH. In this case, when the route via node NH is unavailable or involves a u-turn, as mentioned above, node NB will note that the secondary route from node NB to node ND via node NH has already been tried, and upon replacing the current contents (S) of the virtual source information element 34 with its own node identity, B, it will send the modified message via the primary route via node NC to node ND.

Node NB now forwards the modified Setup Request message to node NC. Upon receipt of the message, node NC first checks whether it is the destination node for that message, and then using the retrieved identity of the virtual source information element 34, accesses its routing table 20C, Figure 6, in respect of the source/destination pair BD, find only a single entry having the address D, and forward the message to the destination node ND.

If, however, there is a fault at node NC, then node NB will have failed to find a route to the destination node ND on both its primary and its secondary routes. There being no further alternative routes in the set B/D, node NB now proceeds to overwrite the current contents (B) of the virtual source information element 34 with the identity of the preceding node NA, A and to change th state of the crankback flag in the crankback information element 36 from reset state to set state. Node NB then sends the modified Setup Request message 30 back to the preceding node NA. Node NA responds to receipt of this modified Setup Request message 30 by then similarly performing the steps of the present invention in accordance with the current contents of the virtual source information element 34, which it will match with its own node identity and proceed on the basis that it is the source of that message.

Thus, node NA does the same thing and it accesses its routing table 20A to find the second entry, A2, for the source/destination pair AD, and retrieves the address, G. The presence of the crankback flag in a set state in the crankback information element 36 causes the nod NA to ignore the entry that has already been used, i.e. the first entry A1, and go straight to the second ntry, A2.

Using the same method, node NG will treat the message as coming from a source A, find the address, D, corresponding to source/destination pair AD, and attempt to route the message to destination node ND.

To sum up, each node has a routing table with three columns, one for the identity of the virtual source node, the second for the identity of the destination node, and the third for the identity of the next node in the route to that destination node. For traffic originating at a node there are always for each destination node at least two entries, the primary and secondary routes, but for transit traffic there is only a single entry for each destination node, i.e. only one of the routes being permitted for use, this being usually, but not always, the primary route from that node to the destination node.

The above described method has following advantages:

(i) it allows loop free routes to be specified for sparsely connected networks under single element, i.e. node or link, failure conditions with only a limited loop prevention mechanism in operation.

(ii) it minimises the operation of crankback under single element failure conditions.

(iii) it can operate successfully with either "crankback to source" or "hop by hop crankback" under failure conditions.

(iv) if used with "hop by hop crankback" it will lead to shorter alternative routes than source routing, but will provide the same resilience advantages as source routing.

(v) it could be used to implement load sharing.

(vi) provided that the source routes are node disjoint, for each source-destination combination, only one routing table entry may be needed at every switch except for the source switches, which always require a set of at least two.

THIS PAGE BLANK (USPTO)

CLAIMS

1. A method of routing in a communications network of interconnected nodes, the nodes being arranged to generate messages, each message having a destination information element containing the identity of a destination node for that message, a source information element containing the identity of the source node of that message, and a virtual source information element initially containing the identity of that source node, the method comprising performing at each node the steps of:

(a) retrieving from a message the contents of its destination information element and its virtual source information element;

(b) comparing the retrieved destination node identity with its own node identity; and, when it is not the destination node for that message,

(c) comparing the retrieved virtual source node identity with its own node identity and,

in the event of a match at step (c),

(d) sending the message to the destination node via the highest ranking available route of a ranked set of alternative routes for that destination node, and

in the event of a mismatch at step (c),

(e) sending the message to the destination node via a predetermined alternative route for that destination node; and

in the event of the unavailability of the predetermined alternative route in step (e),

(f) replacing the content of the virtual source information element of the message with its own node identity and performing step (d), and

if no route is available,

(g) replacing the content of the virtual source information element of the message with the node identity of the node from which the message was received and sending the message back to that node from which it was received.

2. A method as claimed in claim 1, wherein said predetermined alternativ route is one of the set of alternative routes for that destination node.

3. A method as claimed in either claim 1 or claim 2, wherein at least two of the set of alternative routes are node-disjoint routes.

4. A method as claimed in any one of claims 1 to 3, wherein step (g) further comprises changing the state of a flag in a crankback information element of the message, and step (d) comprises ignoring the alternative route that had previously been used for that message.

5. A node for use in a communications network of interconnected nodes, the node being arranged to generate messages, each message having a destination information element containing the identity of a destination node for that message, a source information element containing the identity of the source node of that message, and a virtual source information element initially containing the identity of that source node, and being arranged

to retrieve from a message the contents of its destination information element and its virtual source information element;

to compare the retrieved destination node identity with its own node identity; and, when it is not the destination node for that message,

to compare the retrieved virtual source node identity with its own node identity and,

in the event of a match between the retrieved virtual source node identity and its own node identity,

to send the message to the destination node via the highest ranking available route of a ranked set of alternative routes for that destination node, and

in the event of a mismatch between the retrieved virtual source node identity and its own node identity,

to send the message to the destination node via a predetermined alternative route for that destination node; and

in the event of the unavailability of said predetermined alternative route,

to replace the content of the virtual source information element of the message with its own node identity and to send the message to the

destination node via the highest ranking available route of the ranked set of alternative routes for that destination node, and

if no route is available,

to replace the content of the virtual source information element of the message with the node identity of the node from which the message was received and to send the message back to that node from which it was received.

6. A node as claimed in claim 5, and further arranged to respond to no route being available, by changing the state of a flag in a crankback information element of the message, and to respond to receipt of a message containing a crankback flag in a changed state, by ignoring the alternative route that had previously been used for that message.

7. A communications network comprising interconnected nodes as claimed in either claim 5 or claim 6.

8. A method of routing in a communications network, substantially as hereinbefore described with reference to the drawings.

9. A node for use in a communications network, substantially as hereinbefore described with reference to the drawings.

THIS PAGE BLANK (USPTO)

ABSTRACT

ROUTING IN A COMMUNICATIONS NETWORK

A routing algorithm having particular advantage in sparsely connected networks in which nodes have a ranked set of alternative routes to a destination node, these routes being node-disjoint. Setup messages have an additional information element for the identity of a virtual source node, and a source node inserts its own identity in the virtual source information element. Unless a node is the destination for a message, it examines the content of the virtual source information element of a message, and if there is no match with its own identity it selects from its routing table a predetermined alternative route for the destination node. If that route is unavailable, the node replaces the content of the virtual source information element with its own identity, and performs routing on the basis that there is now a match with its own identity, i.e. it behaves as if it had generated the message. It selects from its routing table the highest ranking of the set for the destination node, and in the event of a fault on the highest ranking route, tries one or more lower ranking routes. If no route is available, the node replaces the content of the virtual source information element with the identity of the node from which it was received, and sends the message back to the node from which it was received.

Figure (2)

THIS PAGE BLANK

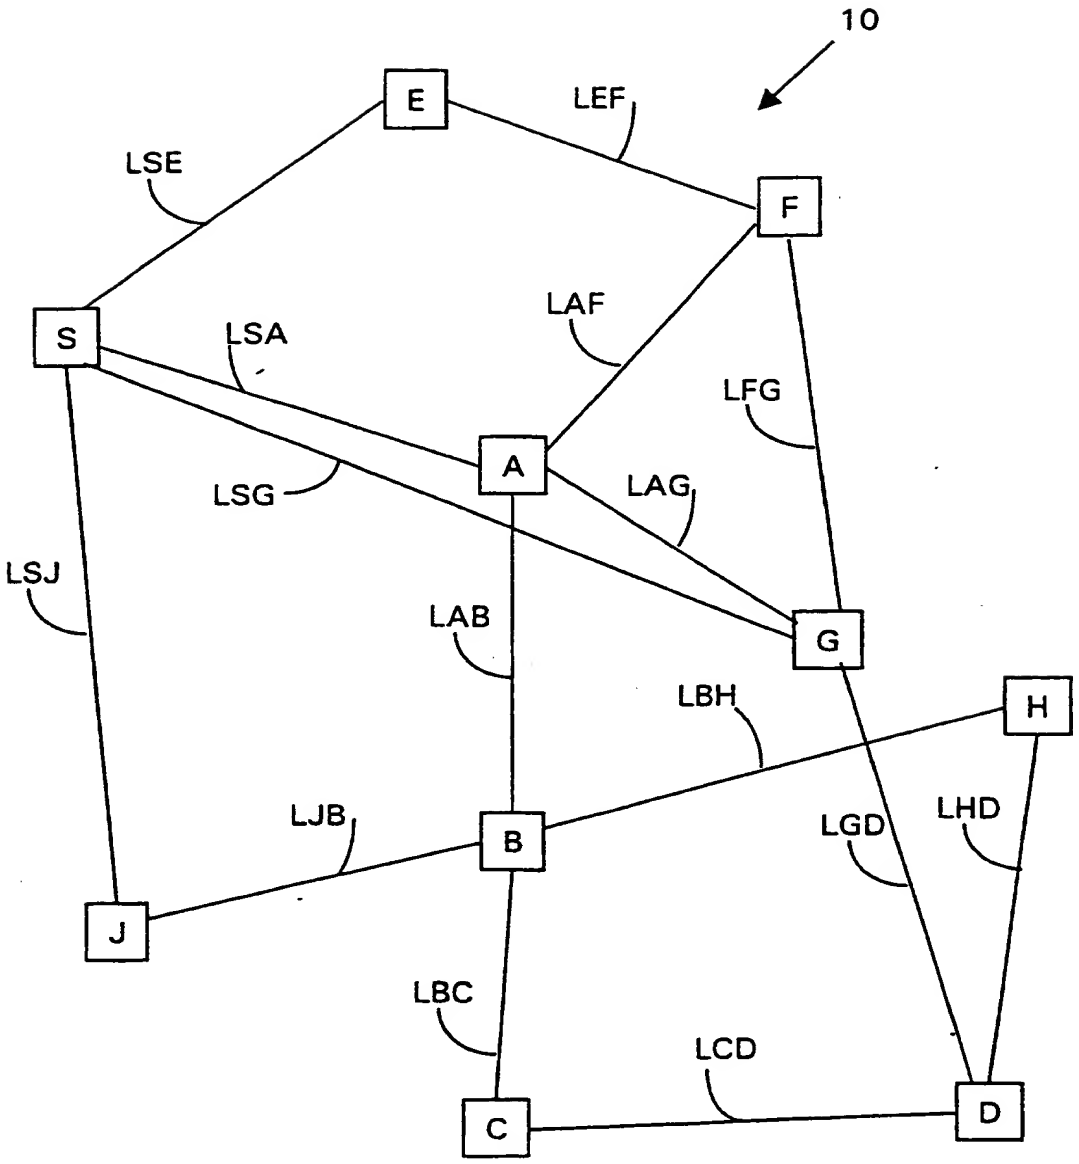
10
E
LEF
LSE
F
S
LSA
LAF
LFG
A
LSG
LAG
LSJ
LAB
G
LBH
H
LJB
B
LGD
LHD
J
LBC
LCD
C
D

Fig. 1

20S

SOURCE	DESTINATION	ADDRESS
S1	D	A
S2	D	E
S1	A	A
S2	A	J
S1	F	A
S2	F	E
S3	F	G
J	F	E
J	G	A
E	B	J
F	J	J

Fig. 3

20A

SOURCE	DESTINATION	ADDRESS
A1	D	B
A2	D	G
S	D	B

Fig. 4

20B

SOURCE	DESTINATION	ADDRESS
B1	D	H
B2	D	C
S	D	H

Fig. 5

20C

SOURCE	DESTINATION	ADDRESS
C1	D	D
C2	D	B
B	D	D

Fig. 6

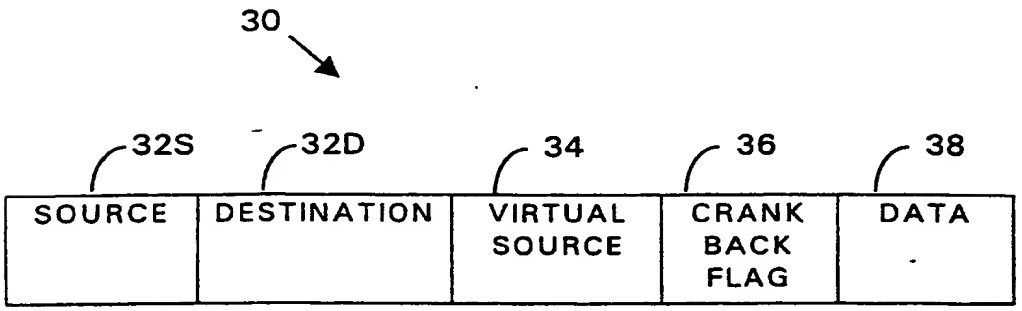

Fig. 2

THIS PAGE BLANK (USPTO)